Adaptive Mesh Refinement for Storm Surge

Kyle T. Mandli[a], Clint N. Dawson[a]

[a]*Institute for Computational Engineering and Science, University of Texas at Austin, 201 E 24th ST. Stop C0200, Austin, TX 78712-1229, USA*

arXiv:1401.5744v1 [math.NA] 22 Jan 2014

Abstract

An approach to utilizing adaptive mesh refinement algorithms for storm surge modeling is proposed. Currently numerical models exist that can resolve the details of coastal regions but are often too costly to be run in an ensemble forecasting framework without significant computing resources. The application of adaptive mesh refinement algorithms substantially lowers the computational cost of a storm surge model run while retaining much of the desired coastal resolution. The approach presented is implemented in the GeoClaw framework and compared to ADCIRC for Hurricane Ike along with observed tide gauge data and the computational cost of each model run.

Keywords: storm surge, Hurricane Ike, adaptive mesh refinement, finite volume methods, shallow water equations

1. Introduction

As computer technology advances, scientists continually attempt to use numerical modeling to better predict a growing number of high-impact geophysical events. In particular, coastal hazards have become an increasing concern as the world's population continues to grow and move towards the coastline, in fact 44% of the world's population lives within 150 km of the coast and 8 of the 10 largest cities in the world lie in that range [1]. As a consequence, loss of life and property is becoming a larger concern than ever before. One of the most recurring and wide spread hazards to many coastal communities is the inundation of coastlines that is associated with strong storms, one part of which is known as storm surge. A storm surge is a rise in the sea accompanying extratropical or tropical cyclones, the strongest examples of which are hurricanes and typhoons. Storm surges can cause massive amounts of damage, as was demonstrated by Hurricane Katrina, which caused an estimated $81 billion of damage [2]. Of the world's largest cities, 4 lie within threat zones from tropical cyclones. With the mounting evidence that severe storms may be increasingly common [3], the task of modeling these events becomes even more critical to communities along the coasts.

Email address: kyle@ices.utexas.edu (Kyle T. Mandli)
URL: http://users.ices.utexas.edu/~kyle (Kyle T. Mandli)

Modeling of storm surges was first carried out by local empirical observations. Unfortunately, for more severe storms such as Katrina, these types of prediction can grossly under-predict storm surge size and effect. By the 1960's, scientists started using computer simulations to predict storm surge but, because these simulations were limited in resolution and size, these models had the same short-comings as the empirically-based models. It was not until recently with increased observational evidence, improved efforts in modeling underlying physical processes, and increases in available computational power that substantial progress has been made simulating large-scale storm surge for use in hazard planning.

The current state-of-the-art numerical models for storm surge simulations rely on single-layer depth-averaged equations for the ocean and make assumptions about the ocean's response to a storm. The National Weather Service (NWS) utilizes a storm surge model called "Sea, Lake and Overland Surges from Hurricanes", or SLOSH, which uses local grids defined for many regions of the United States coastline, to make predictions [4]. These simulations are efficient enough that ensembles of runs can be made quickly for multiple different hurricane paths and intensities. This capability can be critical for effective forecasting due to the uncertainty in the storm forecast. The primary drawback to using the SLOSH model is the limited domain size and extents allowed due to the grid mapping used and formulation of the equations.

Another model currently in use is the Advanced Circulation Model (ADCIRC), a finite element model which has been applied to southern Louisiana [5] and recently to Hurricane Ike [6]. One of the key advantages ADCIRC has it its use of an unstructured grid. Unstructured models allow easy application of variable resolution, especially at the coastline where fine scale features need to be resolved. They can also map to coastlines in a way even a cleverly mapped structured grid cannot. Another advantage of unstructured grids relates to the importance of including entire ocean basins for surge predictions [7, 8]. Unstructured grids can allow the domain of the numerical model to stretch well away from coastlines to include ocean basins while reducing the cost of the model by substantially decreasing resolution in the basin compared to the coastal regions. Unfortunately these models, even with the above advantages, can still be computationally costly and require a large amount of computing resources in order to compute ensemble forecasts without the degradation of their resolution benefits.

In this paper we present an alternative computational framework and methodology to bridge the gap between the numerical cost of the unstructured grid storm surge models and the efficient but unresolved models currently in use at the NWS. The approach leverages adaptive mesh refinement (AMR) algorithms to retain the resolution required to resolve coastal inundation but only when necessary so that ensemble calculations are still feasible. This is accomplished by allowing nested structured grids of variable resolution to vary in time and space thereby capturing the spatial advantages of the unstructured grid approach but only when needed, and therefore decreasing the computational cost substantially. The framework in question, GeoClaw, has successfully been used previously for tsunami modeling where similar computational requirements are present [9].

2. Numerical Approach

The mathematical model for storm surge we will consider uses the classical shallow water equations with the addition of appropriate source terms for bathymetry, bottom friction, wind friction, non-constant surface pressure and Coriolis forcing which can be written as

$$\begin{aligned}
&\frac{\partial}{\partial t}h + \frac{\partial}{\partial x}(hu) + \frac{\partial}{\partial y}(hv) = 0, \\
&\frac{\partial}{\partial t}(hu) + \frac{\partial}{\partial x}\left(hu^2 + \frac{1}{2}gh^2\right) + \frac{\partial}{\partial y}(huv) = \quad fhv - gh\frac{\partial}{\partial x}b + \frac{h}{\rho}\left(-\frac{\partial}{\partial x}P_A + \rho_{\text{air}}C_w|W|W_x - C_f|\vec{u}|u\right) \\
&\frac{\partial}{\partial t}(hv) + \frac{\partial}{\partial x}(huv) + \frac{\partial}{\partial y}\left(hv^2 + \frac{1}{2}gh^2\right) = -fhu - gh\frac{\partial}{\partial y}b + \frac{h}{\rho}\left(-\frac{\partial}{\partial y}P_A + \rho_{\text{air}}C_w|W|W_y - C_f|\vec{u}|v\right)
\end{aligned} \tag{1}$$

where h is the fluid depth, u and v the depth-averaged horizontal velocity components, g the acceleration due to gravity, ρ the density of water, ρ_{air} the density of air, b the bathymetry, f the Coriolis parameter, $W = [W_x, W_y]$ is the wind velocity at 10 meters above the sea surface, C_w the wind friction coefficient, and C_f the bottom friction coefficient. The value of C_w is defined by Garratt's drag formula [10] as

$$C_w = \min(2 \times 10^{-3}, (0.75 + 0.067|W|) \times 10^{-3}) \tag{2}$$

and the value of the friction coefficient C_f is determined using a hybrid Chezy-Manning's n type friction law

$$C_f = \frac{gn^2}{h^{4/3}}\left[1 - \left(\frac{h_{\text{break}}}{h}\right)^{\theta_f}\right]^{\gamma_f/\theta_f} \tag{3}$$

where n is the Manning's n coefficient and $h_{\text{break}} = 2$, $\theta_f = 10$ and $\gamma_f = 4/3$ parameters control the form of the friction law.

The numerical approach proposed to solve (1) falls under a general class of high resolution finite volume methods known as wave-propagation methods, described in detail in [11]. These methods are Godunov-type finite volume methods requiring the specification of a Riemann solver to update each grid cell in the domain. On top of these methods adaptive mesh refinement is employed to allow for variable spatial and temporal resolution as the simulation progresses. These methods have been implemented together in GeoClaw, a package that was originally designed to model tsunamis [12] and other depth-averaged flows [9, 13, 14]. The rest of this section is dedicated to describing the salient points of the AMR approach and how storm surge physics are represented in GeoClaw. A brief review of wave-propagation methods can be found in Appendix A along with a basic outline of the Riemann solver employed in Appendix B.

2.1. Adaptive Mesh Refinement

Adaptive mesh refinement is a core capability of GeoClaw as it allows the resolution of disparate spatial and temporal scales common to geophysical applications such as storm surge and tsunamis. The patch-based AMR approach used in GeoClaw employs a set of

overlapping logically rectangular grids that correspond to one of many levels of refinement. The first of these levels, enumerated starting at $\ell = 1$, contains grids that cover the entire domain at the coarsest resolution. The subsequent levels $\ell \geq 2$ represent progressively finer resolutions by a set of prescribed ratios r^ℓ in time and space such that

$$\Delta x^{(\ell+1)} = \Delta x^{(\ell)}/r_x^{(\ell)}, \quad \Delta y^{(\ell+1)} = \Delta y^{(\ell)}/r_y^{(\ell)}, \quad \text{and} \quad \Delta t^{(\ell+1)} = \Delta t^{(\ell)}/r_t^{(\ell)}.$$

Each subsequent level is properly nested within the union of grids in the next level coarser (see Figure 1 for an example of the structure of these nested grids). With this hierarchy of grids, the evolution of the solution follows as:

1. Evolve the level 1 (coarsest) grids one time step to t^{n+1}.
2. Fill the ghost cells of all level 2 grids by temporal and spatial interpolation.
3. Evolve the level 2 grids the number of time steps determined by $\Delta t^{(2)}$ needed to reach t^{n+1}.
4. Recursively continue to fill in ghost cells, evolving each level ℓ after the coarser level $\ell - 1$ has been evolved until all levels are at time t^{n+1}.
5. Fill in regions where the grids overlap with the best available data using interpolation.
6. Adjust coarse cell values adjacent to finer cells to preserve conservation of mass (see [15] for a discussion on this topic).

The key benefit of adaptive mesh refinement is the ability to change resolution as the simulation progresses. This is done in a process that involves using a local criteria to flag each cell that requires refinement to the next level. The algorithm then clusters the flagged cells into new rectangular patches which attempts to minimize the number of grids created and the number of grid cells unnecessarily refined [16]. After the new grid structure is created, the previous solution values are copied into the new grid cells if no change in resolution was required or an interpolation and averaging is performed to either coarsen or refine the data available from a different level of refinement.

For the shallow water equations there are three primary areas where care must be taken when implementing adaptive mesh refinement. The first involves the interpolation of the solution and bathymetry. When interpolating where the water column is at rest but bathymetry is varying, using the depth h as the interpolation field will lead to a sea surface η that is not at rest and will result in the creation of spurious waves. To avoid this, interpolation is done with the sea surface instead. From this the depth is computed and the resulting momentum interpolated. This process can be perfomed in wet cells while conserving mass and momentum. This is not the case in the near shore where cells may change from their wet (or dry) state. In order to avoid spurious waves being created mass cannot be conserved and is either lost or gained. This is a result of different grid resolutions of bathymetry being used through the computation. For these cases it is often desirable to refine coast lines before a wave arrives such that the loss or gain in mass does not effect the wave itself and overall leads to a negligible increase in mass overall in the simulation. In the end, the following properties are true for the interpolation approach used.

- Mass is conserved except possibly near coast lines.

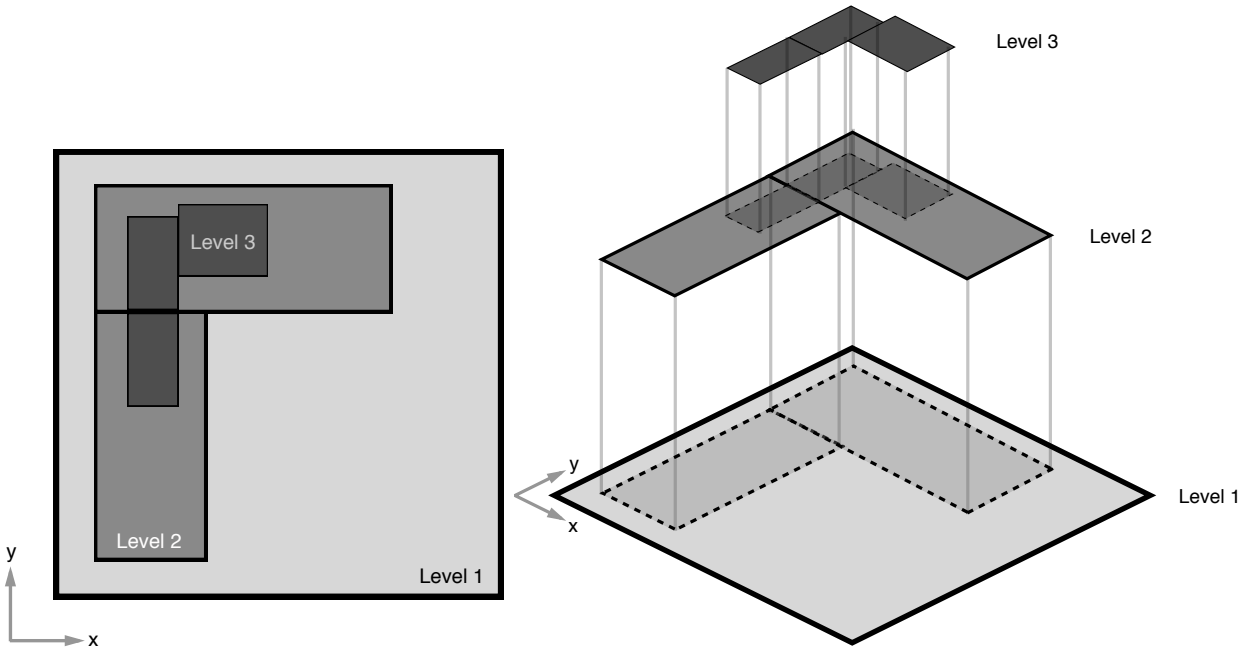

(a) (b)

Figure 1: Views of a sample patch-based AMR grid. Figure 1a shows a planar view with three levels of refinement. In this instance, there are 3 level 3 grids, 2 level 2 grids, and 1 level 1 grid. Figure 1b shows an isometric view of the same grid layout with the z-direction representing different levels.

- Momentum is conserved if mass is conserved or gained and lost if mass is lost.
- New extrema in the sea surface and velocity are not created.

More discussion and the derivation of the rules for coarsening and refinement can be found in [12].

The second area of concern is the form and types of refinement criteria that should be used. In GeoClaw, a number of different solution based criteria are used to determine refinement. The first criteria is triggered if the absolute difference between the initial sea-level $\eta_{\text{sea-level}}$ and the calculated sea-surface η in cell i,j is greater than some tolerance T_{wave}, *i.e.*

$$T_{\text{wave}} < |\eta_{i,j} - \eta_{\text{sea-level}}|$$

will mark the $i^{\text{th}}, j^{\text{th}}$ cell as needing refinement. Additionally water currents can be used as a refinement criteria. This criteria uses a set of speed tolerances specified at each level, only marking a cell as needing refinement if the tolerance T_{speed} for the corresponding level is less than the current speed in that cell. For example, if tolerances for level 1, 2, and 3 were set to 2, 2.5, and 3 m/s respectively, a cell with a speed of 2.6 m/s currently at level 1 then would be marked as needing refinement. Conversely, if the cell is already at level 2 then no refinement is necessary. In addition to these physics based criteria, a user can specify

regions with minimum or maximum refinement constraints. There are also constraints based on initial depth of the water allowing restriction of refinement to only shallower depths. In an instance where a conflict between the physics based criteria and the region constraints occurs the region constraints take precedence.

The last concern of note involves the relative spatial and temporal refinement ratios. Isotropic refinement ratios, when the spatial and temporal refinement ratios are identical, are generally needed in AMR to satisfy the CFL condition. In the case of the shallow water equations and inundation modeling the desired spatial refinement is usually highest at the coastline where the wave speeds are lowest. This allows for the use of anisotropic refinement between the spatial and temporal directions to take advantage of the slower wave speeds located near shore. Along with anisotropic refinement, GeoClaw allows for the automatic determination of the optimal temporal refinement ratios since the wave speed estimate for the shallow water equations is inexpensive to evaluate.

2.1.1. Storm Based AMR Criteria

Additional criteria related to the location and strength of the storm can also be important to resolving storm surge. The first available criteria is based on the distance of a cell from the eye of the storm. As was the case with the speed criteria, these tolerances are specified by level such that a cell will be flagged as needing refinement if it is within the specified radius of the storm's eye, i.e.

$$T_r^\ell < r_{\text{eye}}$$

for level ℓ. The second criteria flags based on storm intensity, wind speed tolerances are used in a similar way. Often the wind speed tolerances and eye distance tolerances give similar but slightly different regions of refinement depending, for instance, on the strength of the storm.

2.2. Source Term Evaluation

Since the storm's impact on the ocean depends only on the momentum source terms for the wind-stress and pressure gradients, it is important that all the momentum source terms are handled carefully. Apart from the storm's impact, the rest of the momentum source terms are bathymetric, friction and Coriolis source terms. The bathymetric source terms are handled in the Riemann solver so that well-balancedness can be maintained along with other useful properties (see Appendix B and [17] for these details). For the remaining source terms, GeoClaw uses an operator-splitting approach to evaluate the remaining source terms. This involves solving two simpler problems, $q_t + A(q)q_z + B(q)q_y = S_1(q)$ and $q_t = S_2(q)$, and combining updates of the vector q in a consistent manner. The method employed uses Godunov splitting which involves solving each simpler system alternately for the full time step Δt. Although this approach is formally only first-order accurate, observed errors due to the splitting do not dominate the overall error in practice (see [11] for a more thorough discussion of these issues).

Starting with the friction and Coriolis, the bottom friction terms can be evaluated as

$$(hu)_t = C_f hu \quad \text{and} \quad (hv)_t = C_f hv$$

and are solved using a backwards Euler method for computing the loss of momentum so that

$$(hu)_{ij}^{n+1} = \frac{(hu)_{ij}^{n}}{1 + (C_f)_{ij}\Delta t} \quad \text{and} \quad (hv)_{ij}^{n+1} = \frac{(hv)_{ij}^{n}}{1 + (C_f)_{ij}\Delta t}$$

where the drag coefficient is computed using the previous time step's state as

$$(C_f)_{ij} = \frac{g n_{ij}^2}{(h_{ij}^n)^{-7/3}} \sqrt{[(hu)_{ij}^n]^2 + [(hv)_{ij}^n]^2} \left[1 - \left(\frac{h_{\text{break}}}{h_{ij}^n}\right)^{\theta_f}\right]^{\gamma_f/\theta_f}.$$

The values h_{break}, θ_f, and γ_f are constant parameters that define the hybridization between the Chezy and Manning's n formulation from (3). The Coriolis terms

$$(hu)_t = -fhv \quad \text{and} \quad (hv)_t = fhu$$

are evaluated using a matrix exponential up to the 4th term in the series such that the update becomes

$$\begin{bmatrix} hu \\ hv \end{bmatrix}_{ij}^{n+1} = \begin{bmatrix} hu \\ hv \end{bmatrix}_{ij}^{n} \cdot \begin{bmatrix} 1 - \frac{1}{2}(f\Delta t)^2 + \frac{1}{24}(f\Delta t)^4 & f\Delta t - \frac{1}{6}(f\Delta t)^3 \\ -f\Delta t + \frac{1}{6}(f\Delta t)^3 & 1 - \frac{1}{2}(f\Delta t)^2 + \frac{1}{24}(f\Delta t)^4 \end{bmatrix}$$

where $f = 2\Omega \sin y$ with $\Omega = 2\pi/8.61642 \times 10^4$ and y the longitudinal coordinate in radians.

2.2.1. Storm Representation and Source Term Evaluation

One of the significant additions to GeoClaw needed to model storm surge is the representation of the storm fields and their source terms. In the forecasting setting that we are targeting, pressure and wind fields of a storm are often evaluated based on the empirically derived Holland model which provides profiles of the wind speed and pressure in a hurricane based on storm characteristics [18]. These profiles are then rotated to produce a two-dimensional, rotationally symmetric field. The wind speed profile $W(r,t)$ takes the form

$$W(r,t) = \sqrt{\left(\frac{r_{\text{w}}}{r}\right)^B W_{\max}^2 e^{(1-(r_{\text{w}}/r)^B)} + \frac{(rf)^2}{4}} - \frac{rf}{2} \tag{4}$$

where r is the radial coordinate centered at the eye of the storm and the parameters r_{w} and $W_{\max}$ are given by the storm forecast and represent the radius of maximum winds, and the maximum wind speed, respectively. The Holland B parameter, an empirical fitting parameter from [18], takes the form

$$B = \frac{\rho_{\text{air}}(W'_{\max})^2 e}{P_n - P_c}$$

where P_n and P_c are the background and central storm atmospheric pressures. The value $W'_{\max}$ is a correction form of $W_{\max}$ that accounts for storm translation and movement in spherical coordinates. The pressure profile is similarly

$$P_A = P_c + (P_n - P_c)e^{-(r_{\text{w}}/r)^B}. \tag{5}$$

The Holland wind field velocities from (4) are converted to a cyclonic wind field via

$$\vec{W} = \begin{bmatrix} -W\sin\theta \\ W\cos\theta \end{bmatrix}$$

where θ is the azimuthal angle with respect to the storm's eye. Furthermore, the storm translation speed is added to take into account the relative motion of the storm and atmosphere. The last modification to the fields is a ramping function, $\mathcal{R}(r)$, applied at a radial distance from the storm to smoothly decrease the wind velocity to zero and the pressure to the background pressure via

$$\mathcal{R}(r) = \frac{1}{2}\left(1 - \tanh\left(\frac{r - R_p}{R_w}\right)\right)$$

where r is the radial distance from the eye of the storm, R_w the width of the ramping function, and R_p the storm parameter recording the radius of the last closed iso-bar. $R(r,t)$ is then applied to the wind and pressure fields as

$$W(r) = W(r)\cdot\mathcal{R}(r) \quad \text{and} \quad P(r) = P_n + (P(r) - P_n)\cdot\mathcal{R}(r).$$

A storm's intensity, speed, and location is allowed to vary in time via forecasted and best-track data at the times specified. GeoClaw then reconstructs the wind and pressure fields using linear interpolation between time points so that the storm forecast evolves continuously rather than only at the time points available. In addition to the intensity of the storm, the velocity and location are linearly interpolated taking into account that the storm is traveling on the surface of a sphere. Once the simulation passes the end of the forecasted data, these parameters are extrapolated in time and kept constant.

Turning now to the evaluation of the source terms due to the wind and pressure fields, we utilize source term splitting as described earlier. The wind source term $\rho_{\text{air}}C_w|W|W$ is evaluated using a single forward Euler time-step with

$$(hu)_{ij}^{n+1} = \frac{h_{ij}^n}{\rho}\rho_{\text{air}}(C_w)_{ij}^n|W_{ij}^n|(W_x)_{ij}^n \quad \text{and} \quad (hv)_{ij}^{n+1} = \frac{h_{ij}^n}{\rho}\rho_{\text{air}}(C_w)_{ij}^n|W_{ij}^n|(W_y)_{ij}^n$$

where the coefficient of drag is defined in (2). The pressure term is also handled via a forward Euler approach such that

$$(hu)_{ij}^{n+1} = \frac{h_{ij}^n}{\rho}\frac{P_{i-1j} - P_{i+1j}}{2\Delta x} \quad \text{and} \quad (hv)_{ij}^{n+1} = \frac{h_{ij}^n}{\rho}\frac{P_{ij-1} - P_{ij+1}}{2\Delta y}.$$

The values of Δx and Δy for the second order-centered finite difference used for the pressure gradient were also converted from the latitude-longitude coordinate system to meters.

3. Comparisons

As a demonstration of the advantages of using adaptive mesh refinement for storm surge, GeoClaw was used to simulate Hurricane Ike. GeoClaw's results were then compared to

gauge data taken during the storm from [19], and to ADCIRC results, previously validated for Hurricane Ike in [6]. The intention was to do the comparison in a forecasting type of scenario and as a consequence some forcing terms and resolution were sacrificed. Each simulation was computed 3 days ahead of the approximate landfall time and run to 18 hours after landfall.

3.1. Hurricane Ike

Hurricane Ike was a storm that caused significant devastation with maximum sustained winds of 54 m/s (10 minutes average) and a minimum pressure of 935 mbar, making landfall in the United States near Galveston, TX on September 13th, 2008 [6]. One of the key features of the U.S. landfall was a prominent forerunner, a rise in water elevation that arrived prior to the arrival of the storm, leading to an increase of the overall surge throughout the storm. The source of this prominent forerunner is thought to be Ekman setup on the Louisiana-Texas shelf [19] as the storm moved parallel to the shelf. This also lead to significant setup in the relatively low lying coastal areas of Louisiana and Texas east of Galveston as the storm winds pushed the surge westward. For a more complete description of the processes of interest during Hurricane Ike see [6].

3.2. ADCIRC Run Description

ADCIRC uses a continuous Galerkin finite element method that solves the modified shallow water equations [5, 20, 21]. It has been validated using hindcasted storm data for a variety of hurricanes, in particular for Louisiana and Texas. A key advantage of ADCIRC over many other available models is the unstructured grids it employs to model storm surge allowing for high levels of resolution near the coasts and less in the Gulf of Mexico and Atlantic Ocean. Even with an unstructured grid, however, the number of nodes and elements required to run a storm surge simulation is massive due to the required domain size and resolution. Because of this, ADCIRC employs multi-process parallelism via the message passing interface (MPI) standard to produce results in a reasonable amount of time. Work has been done to make this parallelization as efficient and scalable (on the order of 10,000 processes) as possible [22].

The ADCIRC simulation compared here is based on a forecasting mode commonly used to model surge along the Texas coastline. The grid used contains 3,331,560 nodes and 6,633,623 elements concentrated along the Texas coastline and identical in that region to the grid used to hindcast Hurricane Ike in [6]. The grid cell sizes range from 30 kilometers to 37 meters with a mean size of 700 meters. Along with the bathymetry, the bottom friction coefficient is allowed to vary at every node of the grid. Additionally, because the simulation was run in forecast mode, a Holland [18] based storm model was used and tides, wave-stress, and riverine input were not included in the computation.

3.3. GeoClaw *Run Description*

In this section we describe the setup for the GeoClaw simulation. The data and settings used for this simulation can be obtained at `http://github.com/clawpack/apps/tree/master/storm_surge/gulf/ike` along with the GeoClaw software itself at `http://www.clawpack.org`.

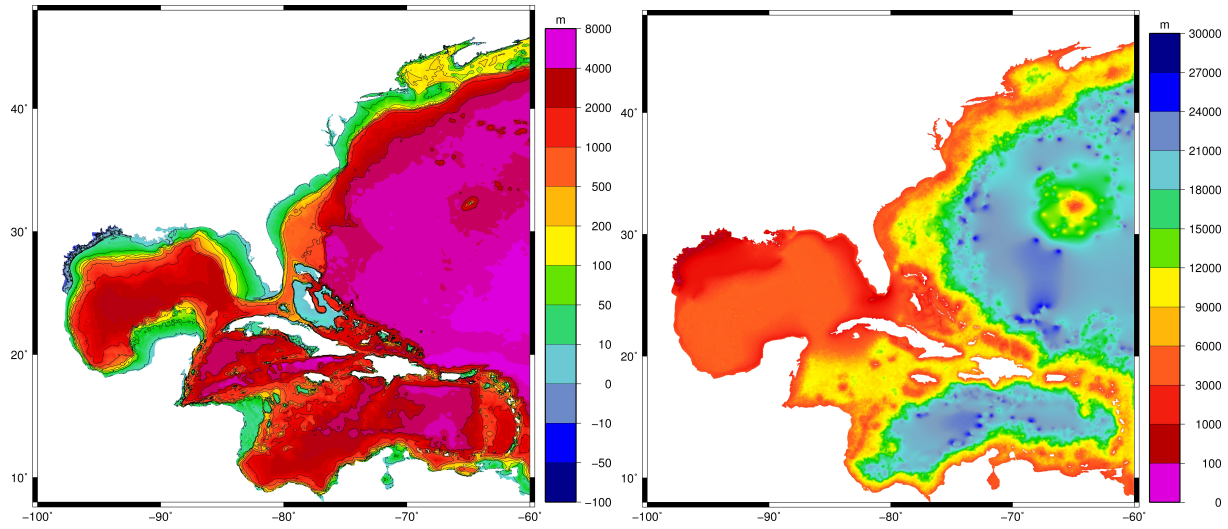

(a) (b)

Figure 2: Characteristics of the grid used in the ADCIRC comparison run. Figure 2a shows the bathymetry used in the ADCIRC grid, Figure 2b shows the resolution of the elements used.

3.3.1. Bathymetry Sources

The bathymetry data used is from the NOAA NGDC US coastal relief model grids with 3 seconds of resolution in the Louisiana-Texas shelf and coastal regions [23]. Outside of these regions, the ETOPO1 global relief bathymetry was used [24]. Sea-level was set to 0.28 meters above the sea-level in the bathymetric sources. This is to account primarily for the swelling of the Gulf of Mexico during the summer along with other effects such as upper layer warming, seasonal riverine discharges, and the measured sea level rise [6].

3.3.2. Variable Friction Field

As was mentioned previously, variable friction can be important to take into account in storm surge simulations. For the GeoClaw results presented, a simple variable friction field was specified based on bathymetry contours relative to sea-level. For regions which were initially above sea-level, a Manning's n coefficient of 0.030 was used and 0.022 for regions initially below sea-level. The one exception to this is in the Louisiana-Texas shelf region (defined as $(25.25°\text{N}, 98°\text{W}) \times (30°\text{N}, 90°\text{W})$ in the simulation) where additional contours at 5 and 200 meters depth were used such that

$$n = \begin{cases} 0.030 & \text{if } b > \eta_{\text{sea-level}} - 10 \\ 0.012 & \text{if } \eta_{\text{sea-level}} - 5 > b > \eta_{\text{sea-level}} - 200 \\ 0.022 & \text{if } \eta_{\text{sea-level}} - 200 > b \end{cases}$$

where $\eta_{\text{sea-level}} = 0.28$ meters as mentioned before. Figure 3b illustrates the resulting values of n through the entire domain. The seemingly counter-intuitive change in friction on the

shelf was suggested in [19] as a means to explain the anomalously large forerunner observed from Ike.

It should be noted that the value of n can change due to refinement. This is currently handled by evaluating the cell depth at the current resolution. For example, in the near shore region a grid cell could be refined to 4 cells which may not have all been above sea-level independently. When this is the case, the depth of each new cell would determine its value of n. If the fields were set based on a given resolution grid (for instance from NASA land use data), a consistent averaging or interpolation is needed.

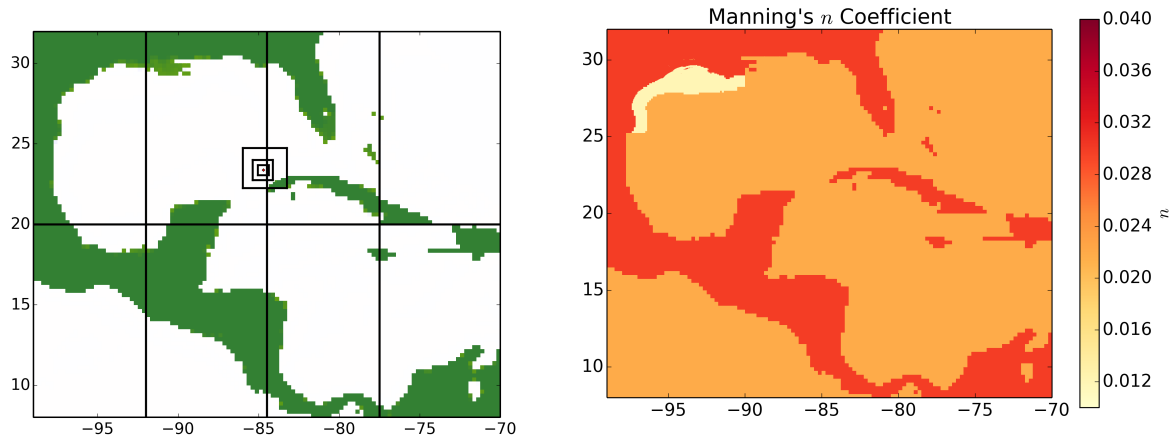

(a) (b)

Figure 3: Figure 3a represents the entire simulation domain including the starting position of the hurricane and refinement patches. Figure 3b shows the Manning's n coefficients used throughout the domain.

3.4. Adaptive Grids

As there was significant setup in the low lying coastal areas east of Galveston, the region of greatest interest, higher refinement was forced in these areas before the storm arrived to capture this effect. The original domain is defined as the region 8° N to 32° N and 99° W to 70° W at 1/4 degree resolution. From this coarsest grid 6 more levels of refinement were used with ratios and resolutions detailed in Table 1. As mentioned in Section 2.1, an optimal refinement ratio in time is used based on the CFL condition on each level. The refinement criteria included sea-level, speed, and storm based criteria and are specified in Table 2. These values were chosen after doing multiple simulations of both Hurricanes Ike and Irene to determine qualitatively the best values in Table 2. Maximum refinement over the entire domain was limited to level 5 away from the coasts of Louisiana and Texas where the full resolution was allowed. As mentioned in Section 2.1, there was an increase in total mass of approximately 0.0015% compared to the original total mass of the domain.

Level	$r^{\ell}_{\Delta x,\Delta y}$	**Resolution (m)**	
		Latitude	Longitude
1		25250	27700
2	2	12600	13850
3	2	6300	6925
4	2	3150	3460
5	6	525	575
6	4	130	144
7	4	32.9	36.1

Table 1: Refinement ratios $r^{\ell}_{\Delta x,\Delta y}$ and effective resolutions at each level. Since the grid is defined in latitude-longitude coordinates, the meters are an approximate resolution in each direction.

Criteria	**Tolerances**
T_{wave}	1 m
T_{speed}	$[1, 2, 3, 4]$ m/s
T_{r}	$[60, 40, 20]$ km
T_{wind}	$[20, 40, 60]$ m/s

Table 2: Refinement criteria tolerances for the sea-surface height T_{wave}, the water speed T_{speed}, the radial distance from the eye T_{r}, and the wind speed T_{wind}. The lists of tolerances correspond to level criteria, i.e. the first entry is the tolerance for moving from level 1 to level 2.

3.5. Simulation Results

Comparisons of the sea-surface and currents produced by GeoClaw and ADCIRC at times before and after land-fall are shown in some of the relevant regions, in particular on the Texas-Louisiana shelf and near Galveston Island and Galveston Bay. For clarity of the comparison the GeoClaw grid patches were not drawn on the plots.

Figure 4 shows the surface displacement in the Texas-Louisiana shelf region produced by GeoClaw and ADCIRC. The simulations match well over most of the shelf and only show differences in regions where the GeoClaw simulation did not fully resolve the features of the surge. Similarly, Figure 5 shows the comparison for the currents in the same region. Again, in the primary surge region, the simulations match well but away from the regions of interest we see additional structure in the ADCIRC simulations, probably due to greater resolution in these areas. These differences may also be due to the more complex frictional field included in the ADCIRC model.

Zooming in on the landfall region, Figures 6 and 7 show surface elevations in the Galveston area before and after landfall respectively. Here, at 18 hours before landfall, we can see that the refinement criteria for GeoClaw has not refined the region significantly. This is

primarily due to the fact that we are plotting from the datum's sea-level rather than the initial sea-level (0.28 meters) used in both simulations. The real difference in sea-level is well below the refinement threshold of 1 meter of surge and therefore refinement has not been triggered. Looking at Figures 8 and 9 that show currents in the Galveston area before and after landfall, we can see that the velocity of the water in the area at 18 hours before landfall is near zero and again refinement is not triggered. Beyond the first snap-shot, we can see that GeoClaw and ADCIRC produce currents that are structurally very similar with GeoClaw producing slightly more diffusive results. Again it should be noted, especially when comparing the current plots, that the ADCIRC simulation has a much more complex friction field than the GeoClaw simulation.

3.5.1. Gauge Data Comparisons

Gauge data was collected during Hurricane Ike off the Texas coast near Galveston [19]. As another point of comparison, this gauge data was compared to both the GeoClaw and ADCIRC numerical gauges. In GeoClaw this data is obtained by nearest neighbor spatial interpolation to the location of the gauge and output in time every time step on the finest grid containing the gauge location. In ADCIRC spatial interpolation is also used and is output at a user defined frequency no smaller than the constant time step taken everywhere. Table 3 describes the location of each gauge and Figures 11, 12, 13, and 14 provide a comparison between the data collected, and the GeoClaw and ADCIRC gauge data.

The first observation that can be made is that the numerical gauges in the GeoClaw simulation are consistently lower than the ADCIRC gauges. This could be due to the resolution near the gauges as the GeoClaw simulation never refines to the highest level of resolution in the gauge locations (as shown by the shading below each gauge figure) where as the ADCIRC simulation's resolution at these gauge locations is approximately . Apart from this, both simulations agree on the timing of the peak surge. Neither the GeoClaw nor the ADCIRC simulations capture the initial peak of surge (the forerunner). This may be due to the use of the Holland based approximated storm fields or model inaccuracies.

Gauge	**Latitude**	**Longitude**	**Figure**
1	29.07 N	95.04 W	11
2	29.28 N	94.71 W	12
3	29.49 N	94.39 W	13
4	29.58 N	94.13 W	14

Table 3: Locations of gauges presented. The background field is of the sea-surface at landfall.

3.5.2. Computational Cost

As one of the objectives of introducing AMR to storm surge forecasting was to reduce the computational cost, we report the overall wall clock time and computational cost of each

of the compared simulations. The ADCIRC simulation was run on 4000 cores of the Stampede computer cluster at the Texas Advanced Computing Center. Stampede is comprised of nodes containing two 8 core Xeon ES-2680 processors with 32GB of memory. These nodes also contain Intel Xeon Phi SE10P co-processors but these were not utilized in the simulation. The GeoClaw simulation was run on a MacBook Pro with an Intel Core i7 chip containing 4 cores and 16GB of memory. As mentioned earlier, ADCIRC utilizes MPI for parallelism while GeoClaw uses OpenMP. Table 4 contains the raw performance and overall computational costs of the ADCIRC and GeoClaw simulations. As another measure of the computational cost in time of the GeoClaw simulation, Figure 15 records the number of grids and grid cells used as a function of time. Although GeoClaw has an overall wall clock time approximately 3.4 times that of the ADCIRC simulation, the computational cost is significantly lessened. It is difficult to fairly compare the equivalent number of degrees of freedom between the GeoClaw and ADCIRC simulations but taking the maximum number of grid cells GeoClaw ever uses as a base, Figure 15 shows the probable source of this decrease in computational cost, the number of grid cells used drastically changes over time.

Package	**Cores Used**	**Wall Clock Time**	**Core Hours**
ADCIRC	4000	35 minutes	2333 hours
GeoClaw	4	2 hours	8 hours

Table 4: Timing and computational cost comparisons between ADCIRC and GeoClaw.

3.6. Discussion

We close this section with some general observations and conclusions regarding the comparison presented. First of all, the GeoClaw simulation compares favorably with the ADCIRC simulation. The peaks surge characteristics are similar although GeoClaw appears more diffusive in this comparison. This may partly be due to the differences between the two simulations in the friction coefficient field but it is more likely to be a combination of insufficient resolution and the nature of the structured grid GeoClaw uses. The ADCIRC grid is designed to mesh waterways and coastlines well where as GeoClaw has to rely on sufficient resolution to represent a channel (unless said waterway is aligned with the chosen grid). It is also interesting to consider Figure 15 that provides some insight into the behavior of the presented simulation. As the storm approaches shore, GeoClaw does not refine much until the storm is approximately 18 hours away from landfall. At this time the resolution increases dramatically and stays relatively high, only decreasing slowly for the rest of the simulation time. An improved refinement criteria may be able to address the apparent discrepancy in resolution while balancing that with decrease of resolution later in the simulation.

Neither simulation fully captures the true surge as was seen in Figures 11, 12, 13, and 14. In a forecasting scenario, capturing the forerunner would be a pertinent feature and it is unclear what caused both simulations to miss this feature. We have also purposely excluded a number of forcing functions (*e.g.* tidal constituents) which may have been important. Currently tides have not been implemented in conjunction with AMR and may be an important

issue to address in the future, especially where tidal forcing is more important such as the northern Atlantic coast.

One important issue that should be mentioned is that the performance of an AMR scheme, in this case GeoClaw, is highly dependent on the type of refinement criteria in conjunction with the base resolution and ratios used. The setup here was picked after some experience running models with GeoClaw but should not be taken as the optimal settings. In the course of writing this article these settings were changed and run multiple times leading to speed-up of a factor of 10 from the initial settings. Clearly more work needs to be spent on what are optimal AMR settings for storm surge modeling.

4. Conclusions

At the outset of this article two capabilities were mentioned as essential for storm surge forecasting, ensemble based calculations and simulations containing resolution sufficient to capture multiple length scales. The numerical models SLOSH and ADCIRC address each of these capabilities independently but not simultaneously. The question that is addressed here then is not if an AMR based code, such as GeoClaw, is better at either of these capabilities separately but if it can satisfy both without overly sacrificing either need. The capability to calculate ensemble based forecasts is answered in Table 4. Given a roughly 2 hour forecasting window, on the same machine running 4 ADCIRC simulations in this window, GeoClaw could be run over 1000 times. It is also clear from the simulations presented that with AMR we are able to capture many of the fine-scale features near coastlines that ADCIRC is able to capture. Combined, these results strongly suggest that AMR is a compelling way to forecast storm surge with both ensemble calculation and high-resolution capabilities.

An ancillary impact of using AMR relates to the amount of effort needed to create the detailed unstructured grids that the ADCIRC modelers actively improve upon. This is an especially important consideration for regions at risk which do not have the type of resources required to create these detailed grids ahead of a storm, both in terms of personnel and software. GeoClaw substantially reduces the effort needed and resources required in order to reach solutions, considering the growing ubiquity of bathymetry and land-use data. Furthermore, GeoClaw also allows the grid to adapt to the storm being run, reducing computational cost.

Future improvements to the GeoClaw framework, as it pertains to storm surge modeling, primarily involve incorporating detailed variable friction fields and investigating possible improvements to the storm field representations. The parallelism present in GeoClaw via OpenMP could be improved. In particular, load-balancing and scalability on high numbers of threads will be critical in the future due to the current trend of higher core counts on consumer chips and the emerging many-core architectures. More broadly, the use of AMR in storm surge forecasting will require additional work on the refinement criteria, balancing resolution with cost. It may also be worthwhile to further delve into what aspect of the AMR implemented here provided the largest impact and adopt them into existing codes which are already in operational use.

Acknowledgments. The authors would like to thank Andrew Kennedy for the gauge observations and the reviewers for their time and effort in editing this article. This research was supported in part by ONR Grant N00014-09-1-0649, the ICES postdoctoral fellowship program, the Gulf of Mexico Research Initiative Center for Advanced Research on Transport of Hydrocarbons in the Environment, and the King Abdullah University of Science and Technology Academic Excellence Alliance.

References

[1] UN Atlas of the Oceans [online].

[2] E. S. Blake, E. N. Rappaport, C. W. Landsea, The Deadliest, Costliest, and Most Intense United States Tropical Cyclones from 1851 to 2006, Tech. Rep. NWS TPC-5, NHC Miami (Apr. 7).

[3] Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, 2007, Tech. rep.

[4] C. P. Jelesnianski, J. Chen, W. A. Shaffer, SLOSH: Sea, Lake and Overland Surges from Hurricanes, Tech. rep., NOAA (1992).

[5] J. J. Westerink, R. A. Luettich, J. C. Feyen, J. H. Atkinson, C. Dawson, H. J. Roberts, M. D. Powell, J. P. Dunion, E. J. Kubatko, H. Pourtaheri, A Basin- to Channel-Scale Unstructured Grid Hurricane Storm Surge Model Applied to Southern Louisiana, Monthly Weather Review (2008) 833.

[6] M. E. Hope, J. J. Westerink, A. B. Kennedy, P. C. Kerr, J. C. Dietrich, C. Dawson, C. J. Bender, J. M. Smith, R. E. Jensen, M. Zijlema, L. H. Holthuijsen, R. A. Luettich, Jr, M. D. Powell, V. J. Cardone, A. T. Cox, H. Pourtaheri, H. J. Roberts, J. H. Atkinson, S. Tanaka, H. J. Westerink, L. G. Westerink, Hindcast and validation of Hurricane Ike (2008) waves, forerunner, and storm surge, Journal of Geophysical Research: Oceans 118 (2013) 1–37.

[7] C. A. Blain, J. J. Westerink, R. A. Luettich, Jr, The influence of domain size on the response characteristics of a hurricane storm surge model, Journal of Geophysical Research 99 (C9) (2007) 467–479.

[8] R. Li, L. Xie, B. Liu, C. Guan, On the sensitivity of hurricane storm surge simulation to domain size, Ocean Modelling 67 (C) (2013) 1–12.

[9] M. J. Berger, D. L. George, R. J. LeVeque, K. T. Mandli, The GeoClaw software for depth-averaged flows with adaptive refinement, Advances in Water Resources 34 (9) (2011) 1195–1206.

[10] J. R. Garratt, Review of Drag Coefficients over Oceans and Continents, Monthly Weather Review 105 (7) (1977) 915–929.

[11] R. LeVeque, Finite Volume Methods for Hyperbolic Problems, Cambridge Texts in Applied Mathematics, Cambridge University Press, Cambridge, UK, 2002.

[12] R. J. LeVeque, D. L. George, M. J. Berger, Tsunami Propagation and inundation with adaptively refined finite volume methods, Acta Numerica (2011) 211–289.

[13] D. L. George, R. M. Iverson, A two-phase debris-flow model that includes coupled evolution of volume fractions, granular dilatancy, and pore-fluid pressure, Debris-flow Hazards Mitigation, Mechanics, Prediction, and Assessment (2011) 367–374.

[14] K. T. Mandli, Finite Volume Methods for the Multilayer Shallow Water Equations with Applications to Storm Surges, Ph.D. thesis, University of Washington (Jun. 2011).

[15] M. J. Berger, R. J. LeVeque, Adaptive Mesh Refinement using Wave-Propagation Algorithms for Hyperbolic Systems, SIAM Journal on Numerical Analysis 35 (1998) 2298–2316.

[16] M. Berger, I. Rigoutsos, An algorithm for point clustering and grid generation, Systems, Man and Cybernetics, IEEE Transactions on 21 (5) (1991) 1278–1286.

[17] D. L. George, Augmented Riemann solvers for the shallow water equations over variable topography with steady states and inundation, Journal of Computational Physics 227 (6) (2008) 3089–3113.

[18] G. J. Holland, An Analytic Model of the Wind and Pressure Profiles in Hurricanes, Monthly Weather Review 108 (1980) 1212–1218.

[19] A. B. Kennedy, U. Gravois, B. C. Zachry, J. J. Westerink, M. E. Hope, J. C. Dietrich, M. D. Powell, A. T. Cox, R. A. Luettich, Jr., R. G. Dean, Origin of the Hurricane Ike forerunner surge, Geophysical Research Letters 38 (8) (2011) L08608.

[20] R. A. Luettich, Jr., J. J. Westerink, Formulation and Numerical Implementation of the 2D/3D ADCIRC Finite Element Model Version 44.XX, Tech. rep. (Jun. 2004).

[21] C. Dawson, J. J. Westerink, J. C. Feyen, D. Pothina, Continuous, discontinuous and coupled discontinuous–continuous Galerkin finite element methods for the shallow water equations, International Journal for Numerical Methods in Fluids 52 (1) (2006) 63–88.

[22] S. Tanaka, S. Bunya, J. J. Westerink, C. Dawson, R. A. Luettich, Jr, Scalability of an unstructured grid continuous Galerkin based hurricane storm surge model, Journal of Scientific Computing 46 (3) (2011) 329–358.

[23] U.S. Coastal Relief Model, NOAA National Geophysical Data Center.

[24] C. Amante, B. W. Eakins, ETOPO1 1 Arc-Minute Global Relief Model: Procedures, Data Sources and Analysis, Tech. Rep. NGDC-24, NOAA (Mar. 2009).

Appendix A. Wave Propagation Methods

Wave-propagation methods start with the finite volume discretization of the domain into cells $\mathcal{C}_i$ and the objective to evolve the grid cell average defined as $Q_i^n \equiv \frac{1}{\Delta x_i}\int_{\mathcal{C}_i} q(x,t^n)dx$. The goal then is to evolve these cell averages forward in time as the system of PDEs $q_t + f(q)_x = 0$ prescribes. To accomplish this, the grid cell averages are used to reconstruct a piece-wise constant representation of the function q throughout the domain. This reconstruction leads naturally to the formulation of Riemann problems at each grid cell interface. A Riemann problem consists of the original system of PDEs on an infinite domain with a jump-discontinuity located at $x = 0$. The solution to a Riemann problem generally consists of m waves (often m is equal to the number of equations in the system of PDEs but this is not a requirement) denoted by $\mathcal{W}^p \in \mathbb{R}^m$ propagating away from the location of the jump-discontinuity traveling at speeds s^p. These waves are related to the original jump-discontinuity via

$$Q_i - Q_{i-1} = \sum_{p=1}^{m} \mathcal{W}_{i-1/2}^p. \tag{A.1}$$

From these waves, a first-order upwind method can be written as

$$Q_i^{n+1} = Q_i^n - \frac{\Delta t}{\Delta x}\left[\mathcal{A}^+\Delta Q_{i-1/2} - \mathcal{A}^-\Delta Q_{i+1/2}\right]$$

where $\mathcal{A}^+\Delta Q$ and $\mathcal{A}^-\Delta Q$ represent fluctuations coming from the right and left cell interfaces respectively and can be defined in terms of the waves as

$$\mathcal{A}^\pm\Delta Q_{i-1/2} = \sum_{p=1}^{m}(s_{i-1/2}^p)^\pm \mathcal{W}_{i-1/2}^p$$

where $s_{i-1/2}^+ = \max(s_{i-1/2}^p, 0)$ and $s_{i-1/2}^- = \min(s_{i-1/2}^p, 0)$. Extensions to 2^{nd}-order are possible using limiters applied to each wave such that the update becomes

$$Q_i^{n+1} = Q_i^n - \frac{\Delta t}{\Delta x}(\mathcal{A}^-\Delta Q_{i+1/2} + \mathcal{A}^+\Delta Q_{i-1/2}) - \frac{\Delta t}{\Delta x}(\tilde{F}_{i+1/2} - \tilde{F}_{i-1/2})$$

where

$$\tilde{F}_{i-1/2} = \frac{1}{2}\sum_{p=1}^{m}|s_{i-1/2}^p|\left(1 - \frac{\Delta t}{\Delta x}|s_{i-1/2}^p|\right)\widetilde{\mathcal{W}}_{i-1/2}^p$$

and $\widetilde{\mathcal{W}}_{i-1/2}^p$ are limited versions of $\mathcal{W}_{i-1/2}^p$.

In the case where the system being considered is linear, the waves $\mathcal{W}^p$ can be written as scalar multiples α^p of the right eigenvectors r^p of the flux where $f(q)_x = Aq_x$ such that

$$\mathcal{W}_{i-1/2}^p = (\ell_{i-1/2}^p)^T(Q_i - Q_{i-1})r_{i-1/2}^p$$

where ℓ^p are the left eigenvectors of A. If the system is non-linear often local linearizations are employed to find an approximate flux Jacobian $\hat{A}_{i-1/2}$ whose eigenvectors are again used to find $\mathcal{W}^p$.

An alternative formulation to the splitting in (A.1) is to instead split the jump in the fluxes at each cell interface as

$$f(Q_i) - f(Q_{i-1}) = \sum_{p=1}^{m} \mathcal{Z}^p_{i-1/2}$$

where now $\mathcal{Z}^p_{i-1/2}$ waves carrying a jump in the fluxes. This formulation, called the f-wave approach, has the advantage that spatially dependent flux functions and source terms can be incorporated directly into the Riemann solver. In the case of an f-wave splitting, the fluctuations $\mathcal{A}^+\Delta Q$ and $\mathcal{A}^-\Delta Q$ can be written as

$$\mathcal{A}^{\pm}\Delta Q_{i-1/2} = \sum_{p=1}^{m} \text{sgn}(s^p_{i-1/2})\mathcal{Z}^p_{i-1/2}.$$

The accuracy again can be improved as before by limiting the f-waves.

Appendix B. Augmented Riemann Solver

The augmented Riemann solver used in GeoClaw was originally proposed in [17] where the novel approach of splitting both the conserved quantities Q and the fluxes $f(Q)$ into waves was explored. Here the augmented set of waves take the form

$$\begin{bmatrix} Q_i - Q_{i-1} \\ f(Q_i) - f(Q_{i-1}) \end{bmatrix} = \begin{bmatrix} \sum_{p=1}^{m} \mathcal{W}^p_{i-1/2} \\ \sum_{p=1}^{m} \mathcal{Z}^p_{i-1/2} \end{bmatrix} \tag{B.1}$$

leading to $2m$ waves. Since only m waves are needed to update the m equations, these extra waves can be used to provide many desirable properties to the Riemann solution.

Another primary feature of the Riemann solver in GeoClaw is the direct inclusion of the bathymetry source term $-ghb_x$ into the Riemann solver. One important consequence of this incorporation in conjunction with the augmented approach is the genesis of a "steady state" wave. As was shown in [17] this additional wave preserves steady states where $(hu)_x = 0$, a much more general class of steady states than is usually considered. Additionally this extra wave can be used to better represent large rarefaction waves which are commonly found at the wet-dry front.

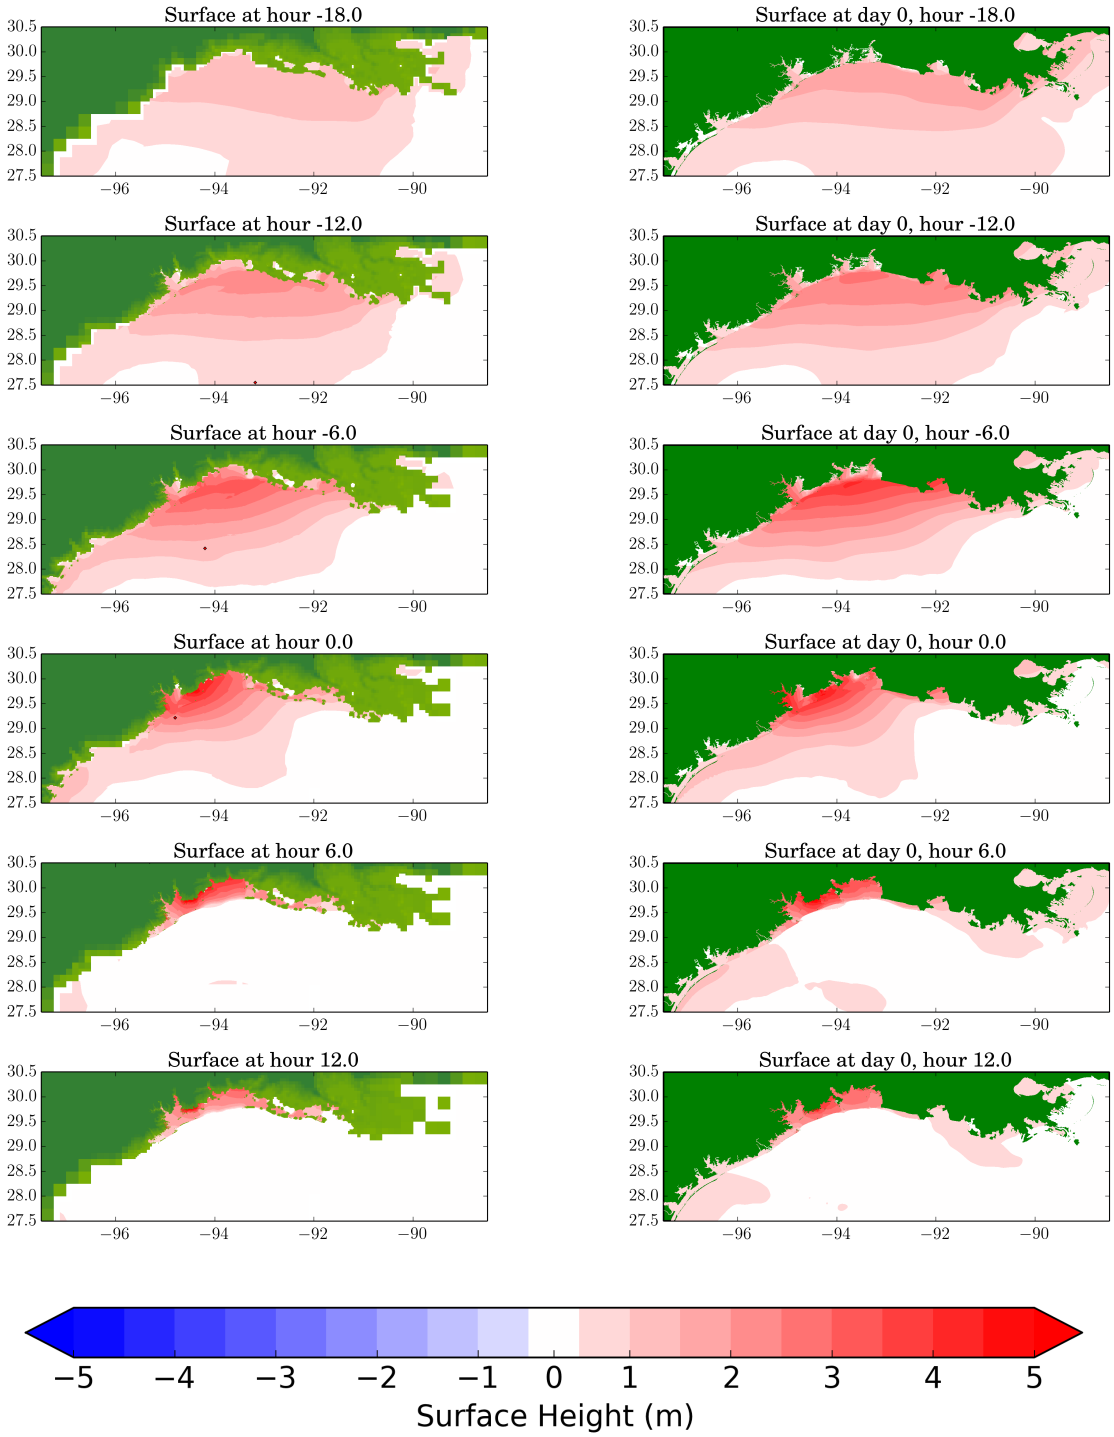

Figure 4: Surface elevations on the Texas-Louisiana shelf produced by GeoClaw (on the left) and ADCIRC (on the right) before and after Hurricane Ike makes landfall.

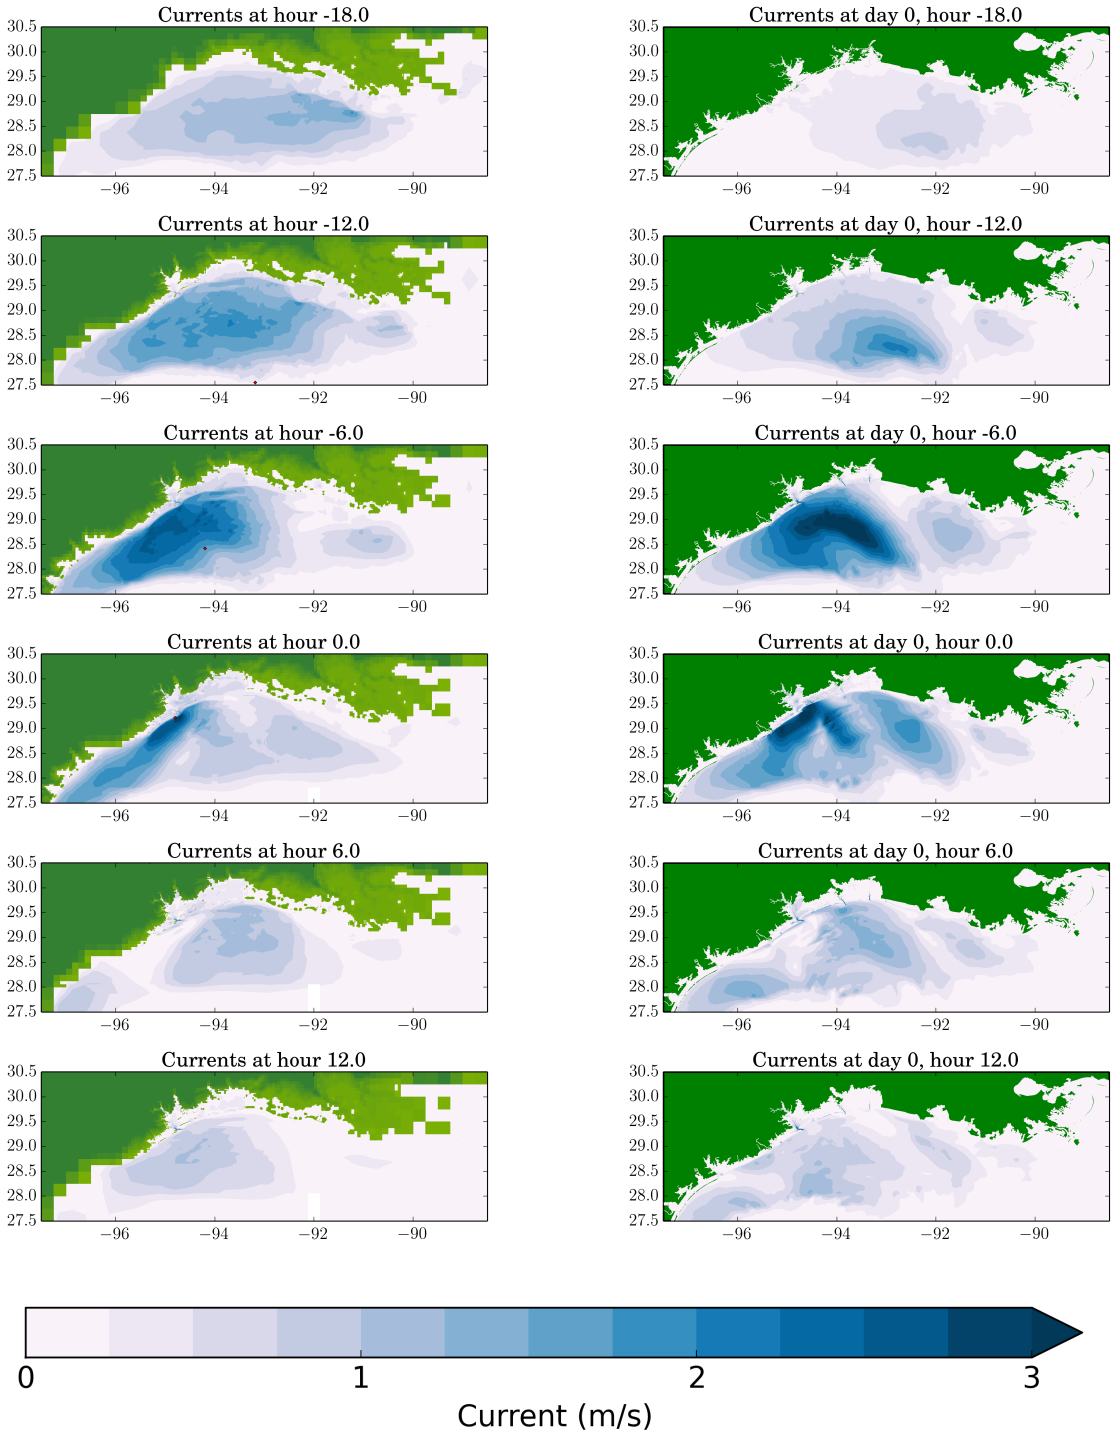

Figure 5: Currents on the Texas-Louisiana shelf produced by GeoClaw (on the left) and ADCIRC (on the right) before and after Hurricane Ike makes landfall.

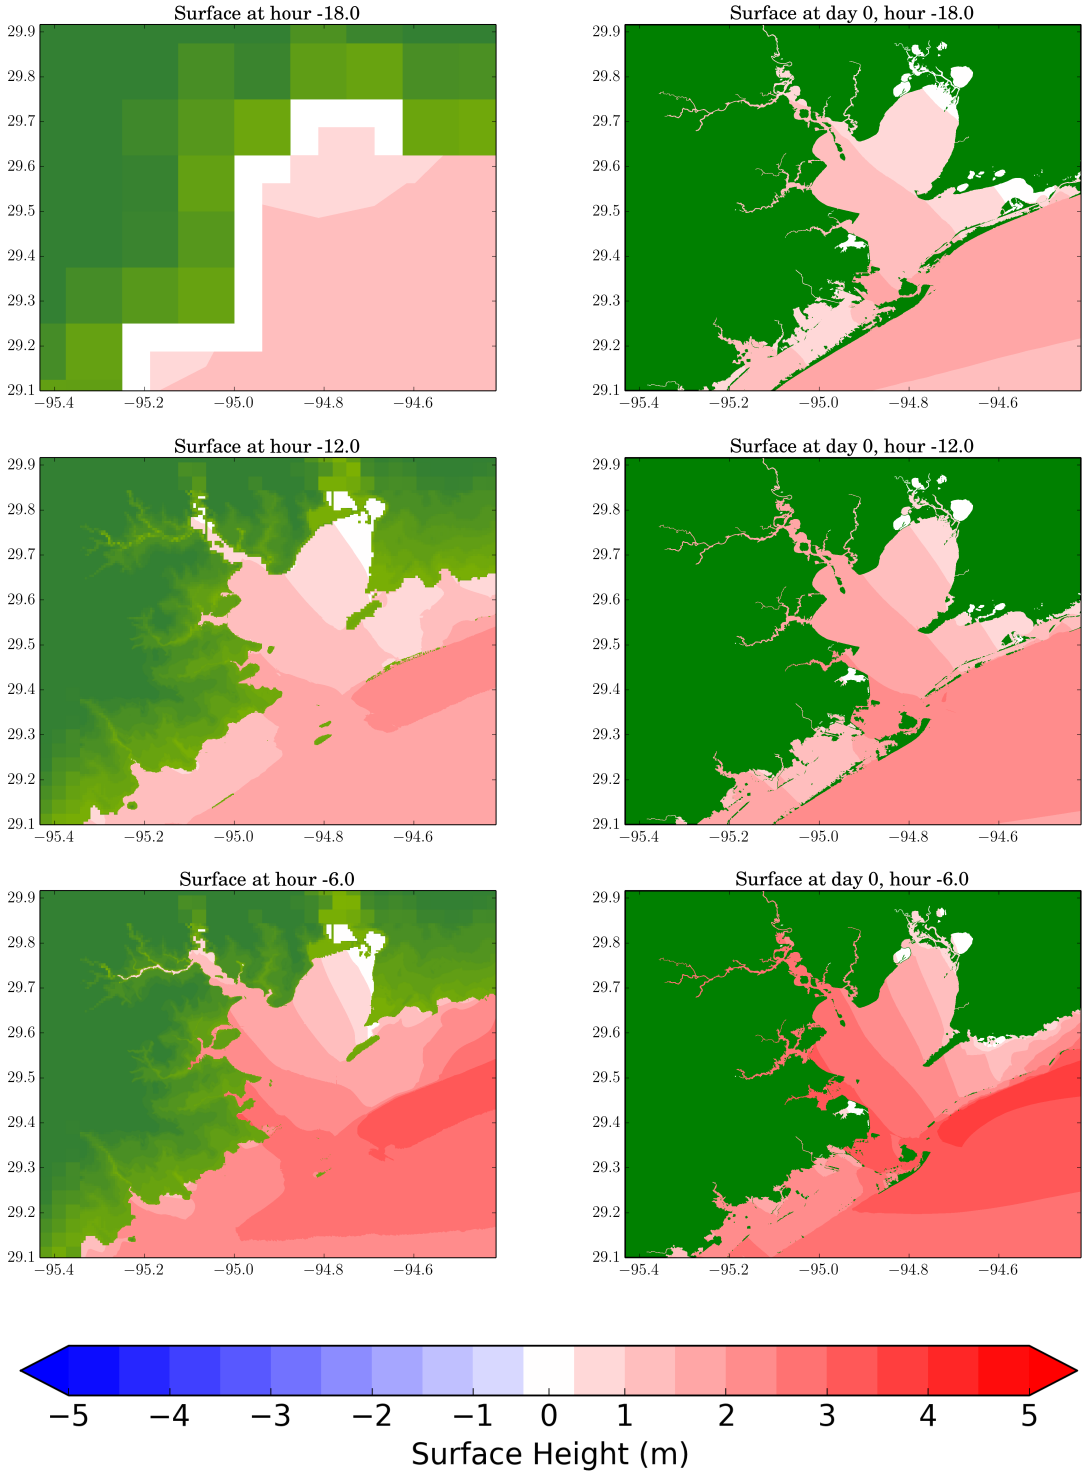

Figure 6: Surface elevations in the Galveston Bay region produced by GeoClaw (on the left) and ADCIRC (on the right) before Hurricane Ike made landfall.

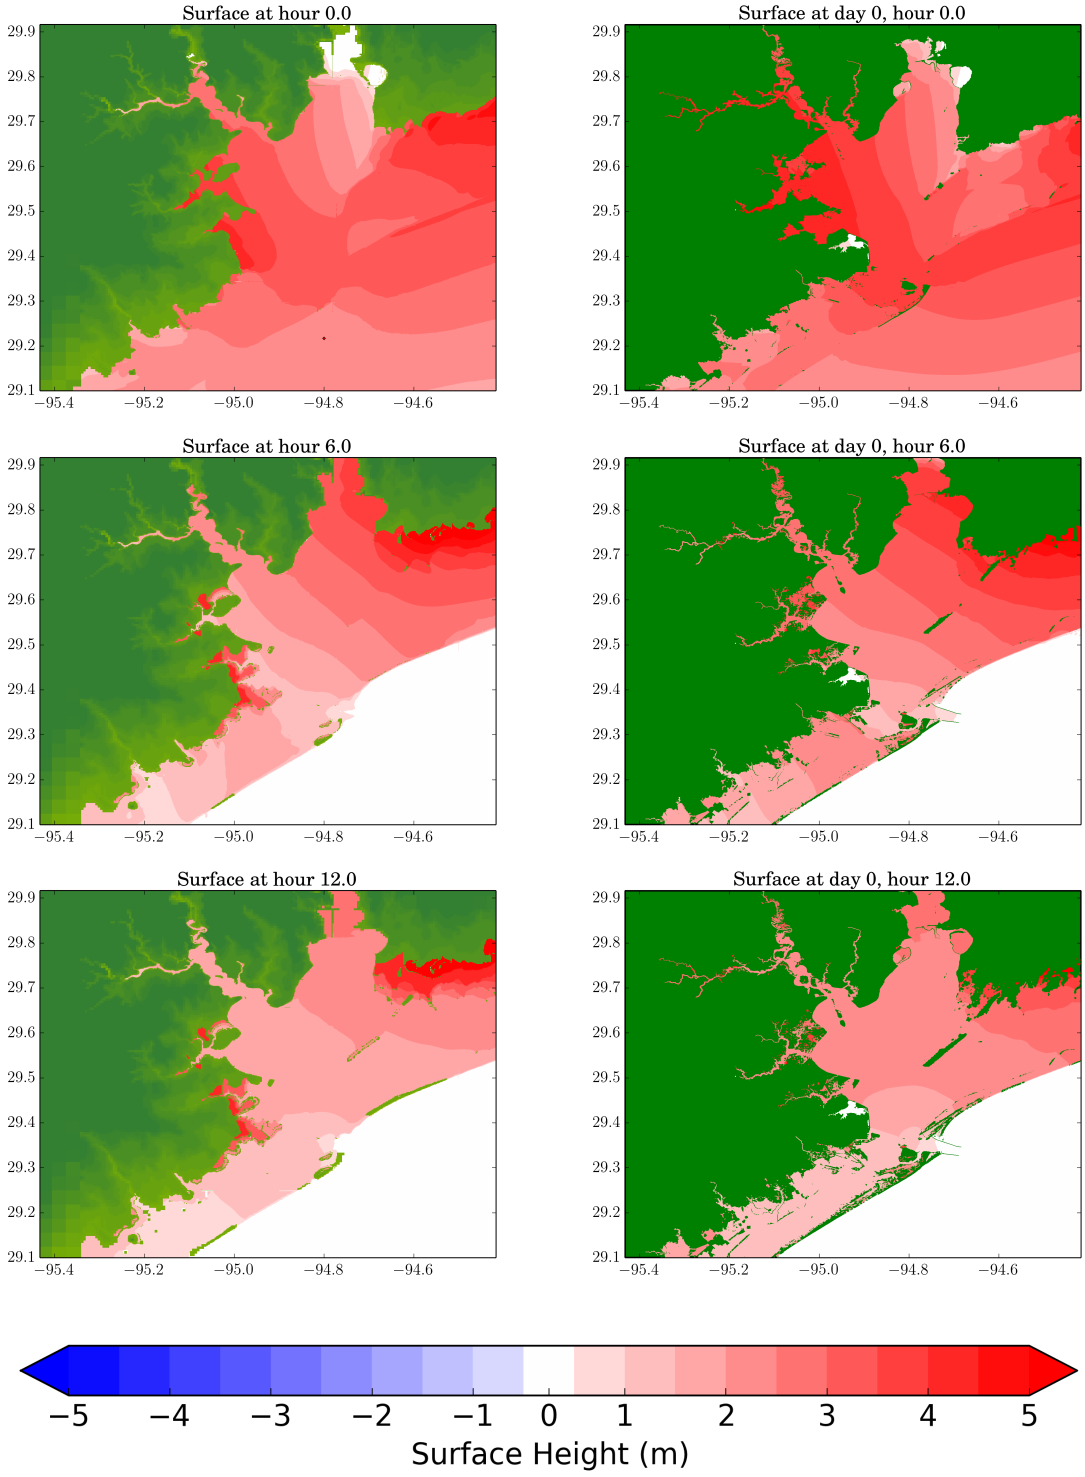

Figure 7: Surface elevations in the Galveston Bay region produced by GeoClaw (on the left) and ADCIRC (on the right) after Hurricane Ike made landfall.

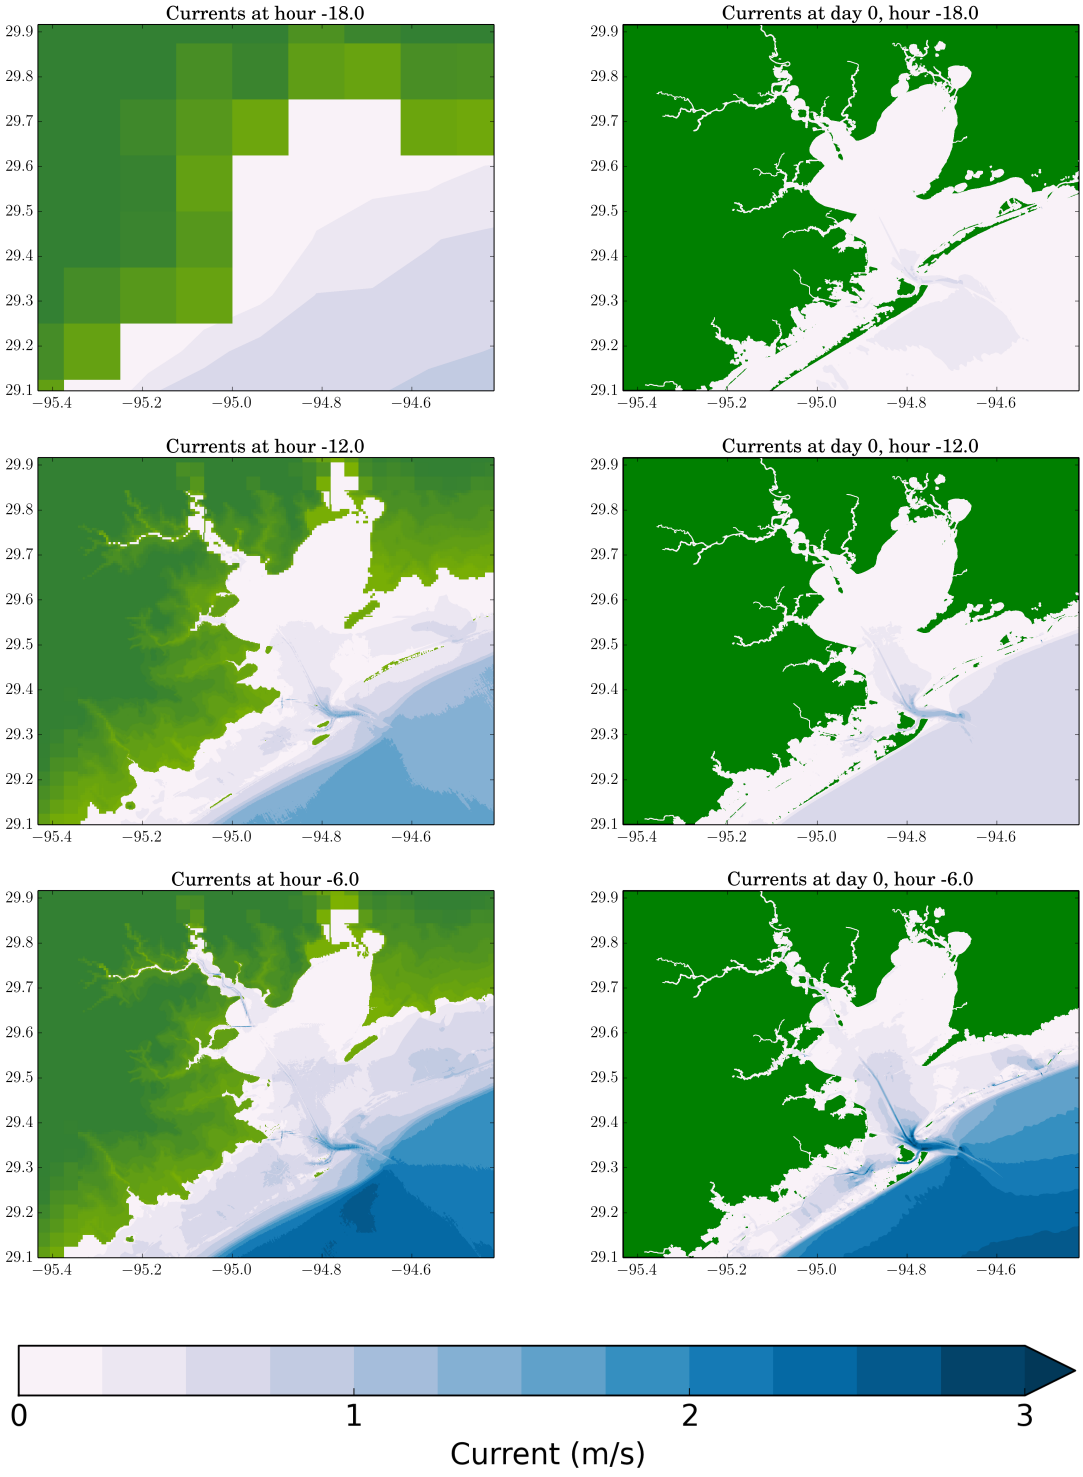

Figure 8: Currents in the Galveston Bay region produced by GeoClaw (on the left) and ADCIRC (on the right) before Hurricane Ike makes landfall.

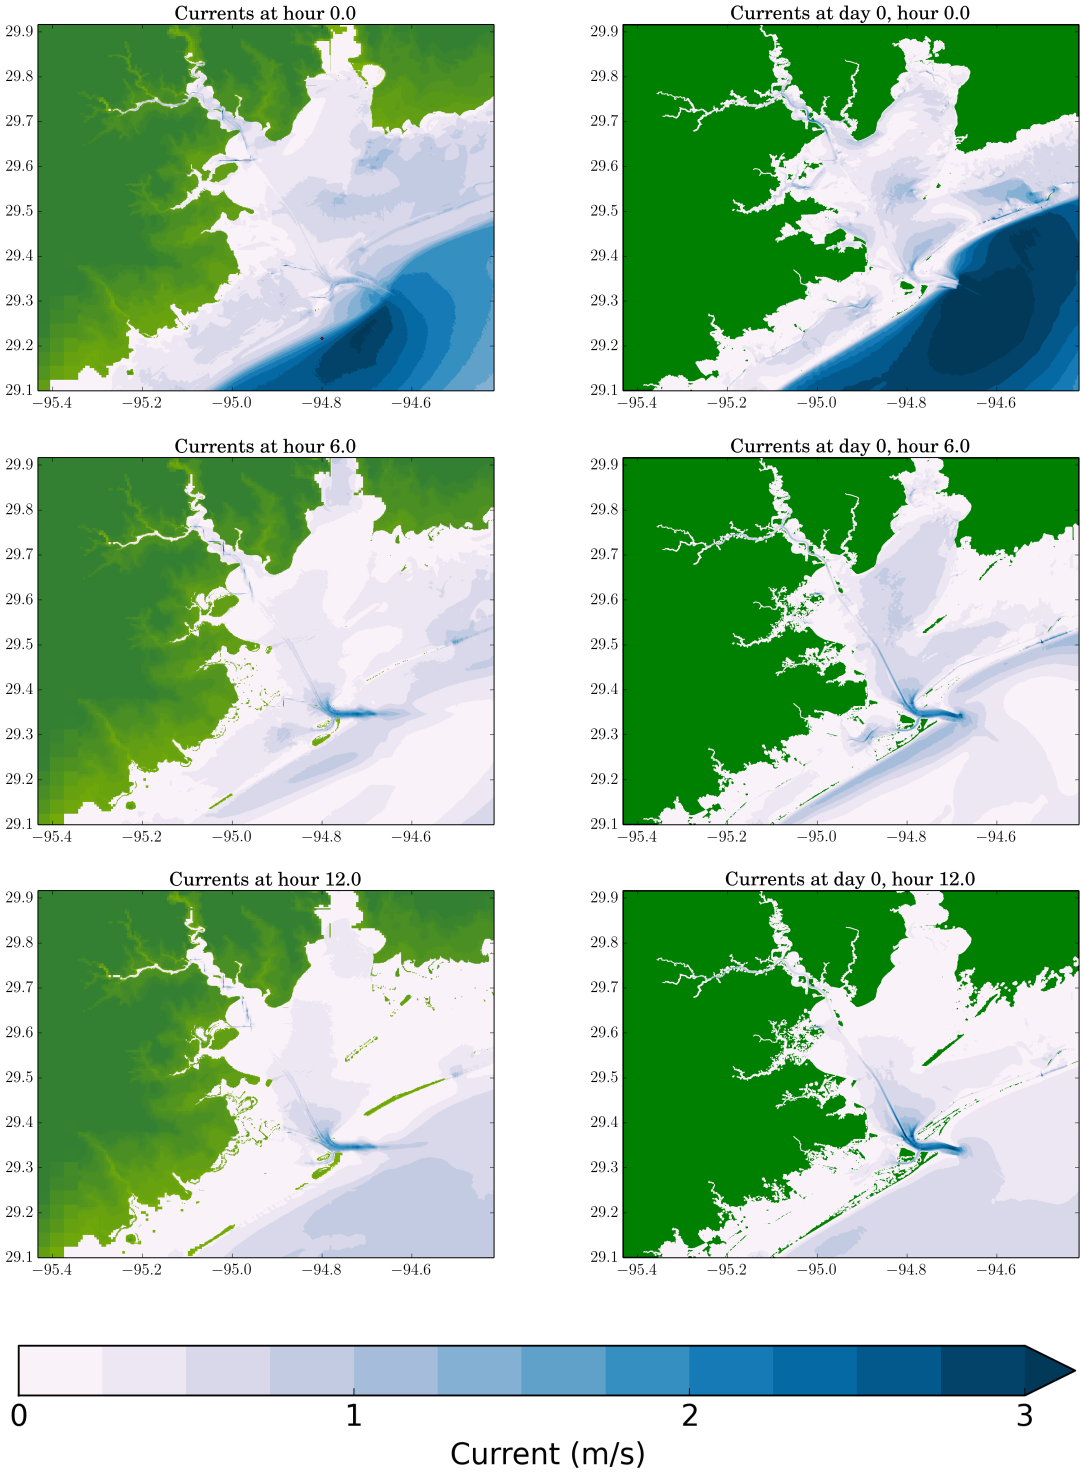

Figure 9: Currents in the Galveston Bay region produced by GeoClaw (on the left) and ADCIRC (on the right) after Hurricane Ike makes landfall.

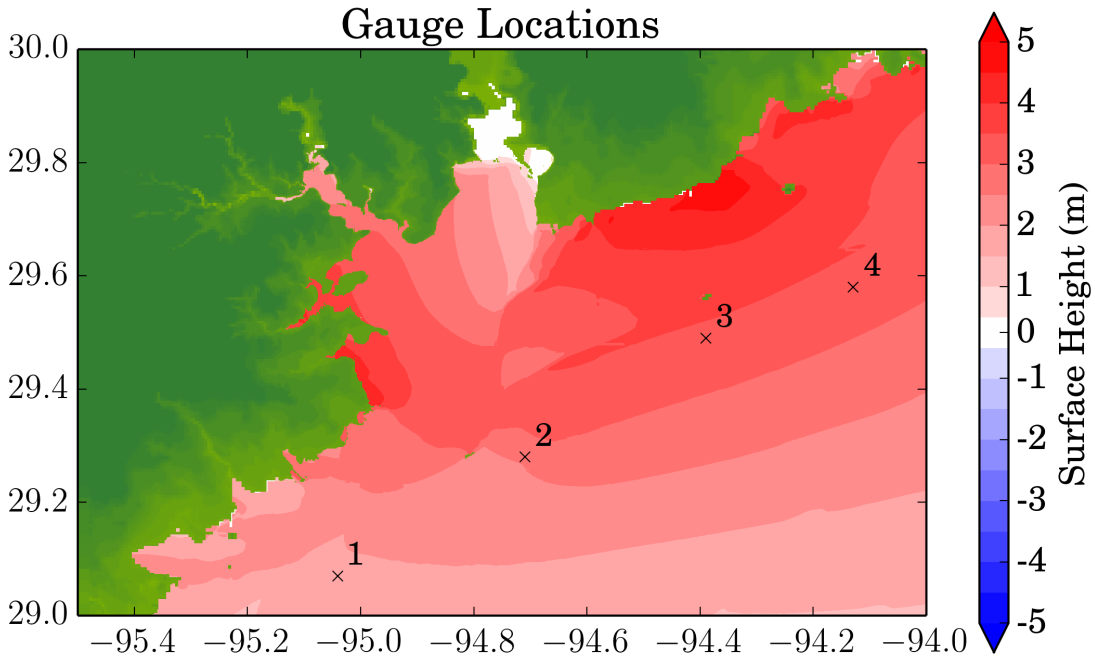

Figure 10: Locations of gauges presented. The color is indicative of the surge heights at landfall (see Figure 7) to give a perspective on where the gauges are relative to the surge.

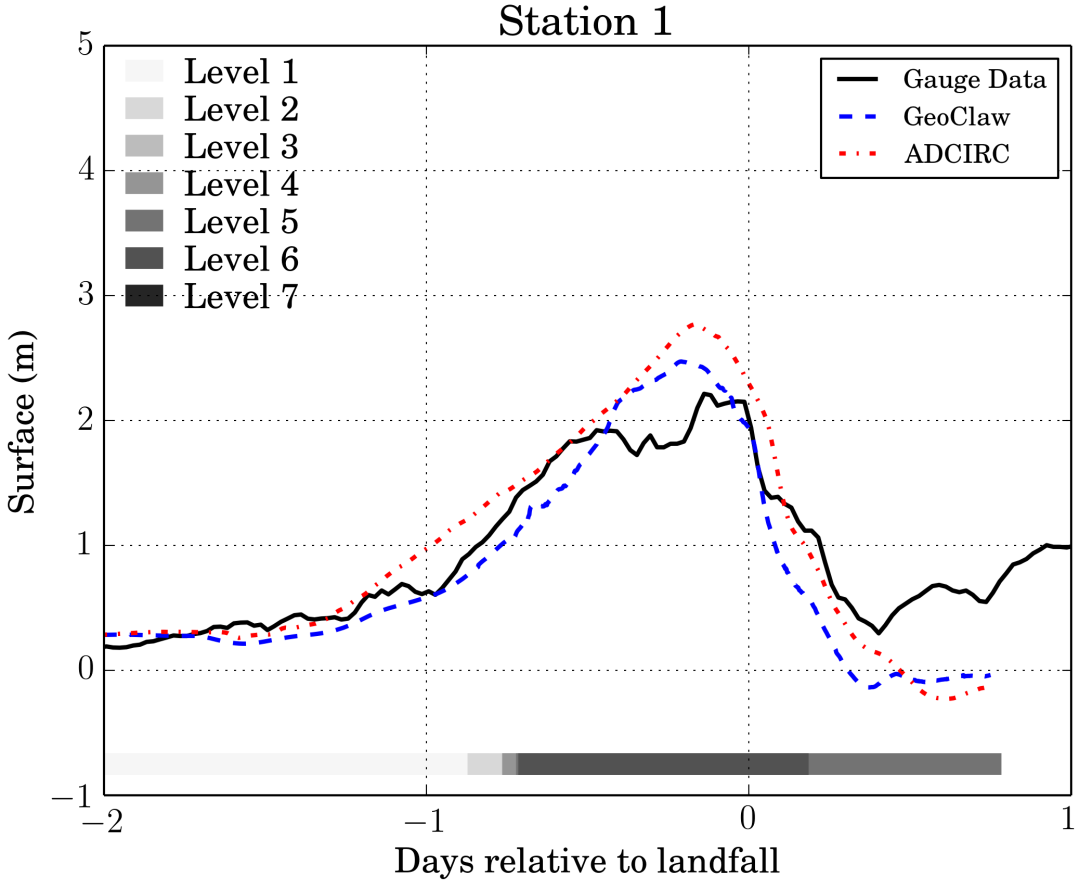

Figure 11: Comparison of numerical gauge data from GeoClaw and ADCIRC and collected data during Hurricane Ike at location 1 from Table 3. The shading below the gauge data represents the refinement level used to record the gauge.

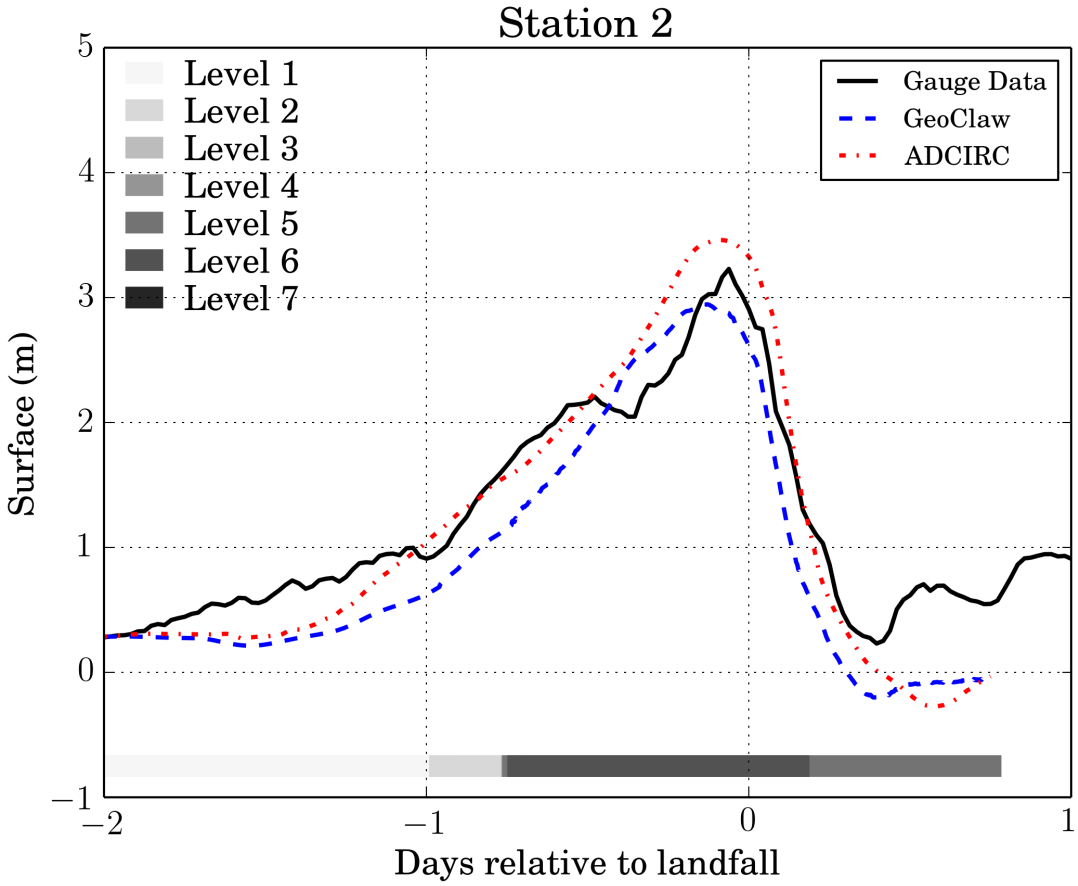

Figure 12: Comparison of numerical gauge data from GeoClaw and ADCIRC and collected data during Hurricane Ike at location 2 from Table 3. The shading below the gauge data represents the refinement level used to record the gauge.

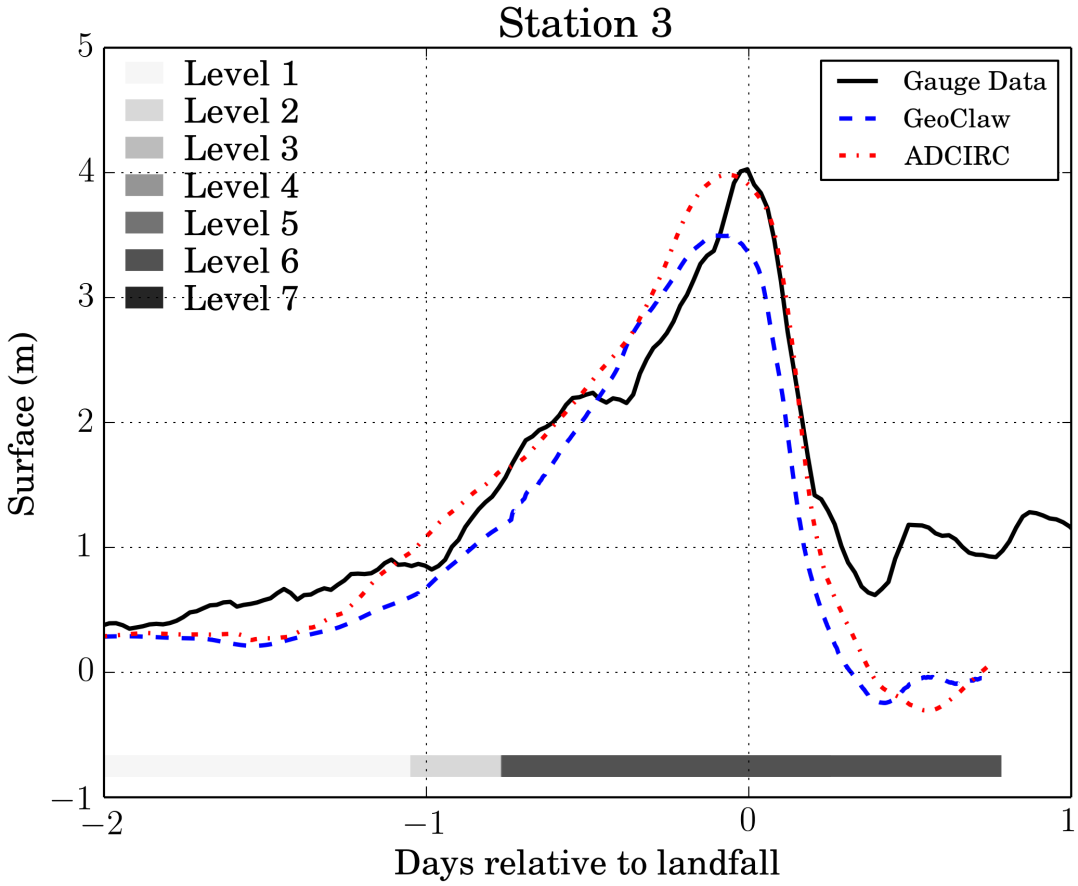

Figure 13: Comparison of numerical gauge data from GeoClaw and ADCIRC and collected data during Hurricane Ike at location 3 from Table 3. The shading below the gauge data represents the refinement level used to record the gauge.

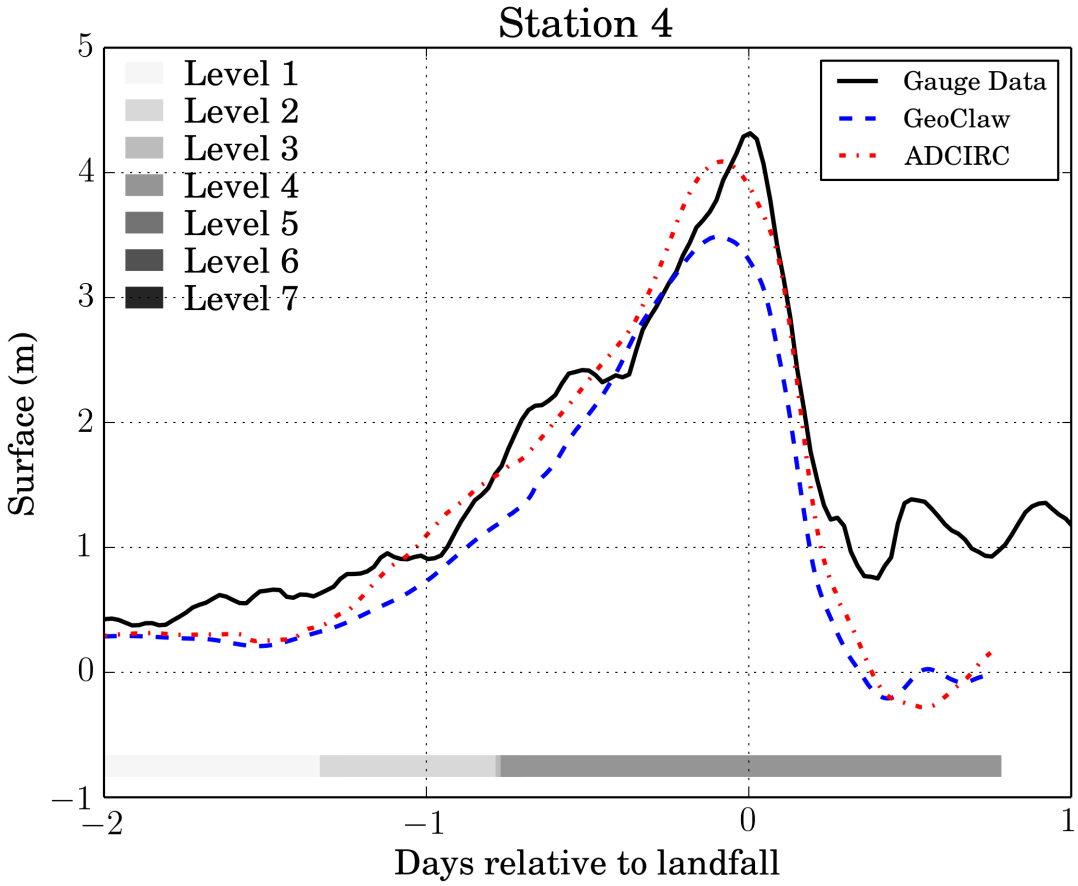

Figure 14: Comparison of numerical gauge data from GeoClaw and ADCIRC and collected data during Hurricane Ike at location 4 from Table 3. The shading below the gauge data represents the refinement level used to record the gauge.

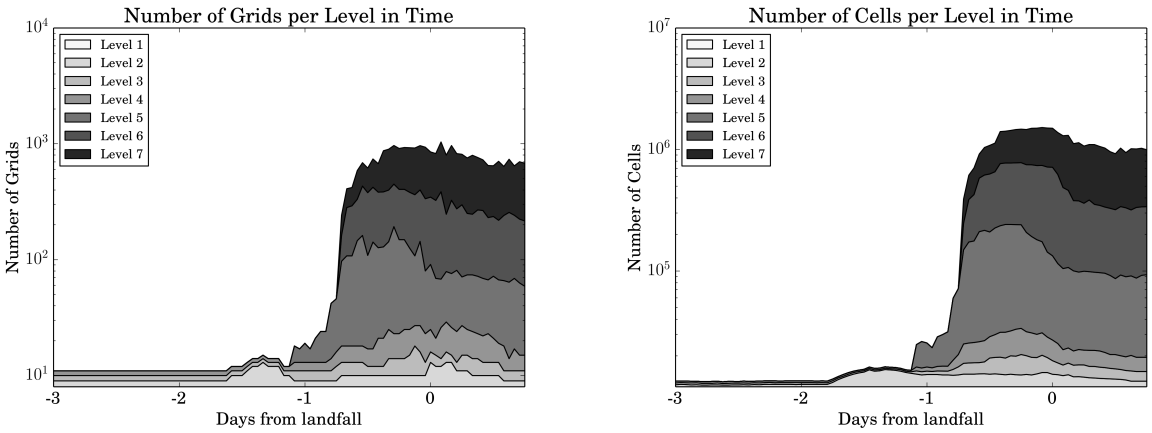

(a) (b)

Figure 15: Domain resolution statistics for the GeoClaw simulation. In Figure 15a a stack plot is shown representing the total number of grids per level and total. Similarly in Figure 15b the number of grids cells is shown.